Stress hedging in portfolio construction

Scenario stress testing is a useful and increasingly popular approach to assess portfolio performance under different market conditions. In this article, Mehmet Bilgili, Maurizio Ferconi and Alex Ulitsky focus on how to directly incorporate stress scenario information into portfolio construction as an additional constraint to control for potential losses and risks. To broaden the applicability of stress testing, the authors propose a robust, constrained optimisation approach to handle uncertainty in scenario parameters. An oil crisis event is used as a numerical illustration

Recent event-driven fluctuations in financial markets have made investors increasingly concerned with finding ways to assess portfolio performance under different economic conditions. A growing number of financial practitioners and academics have begun to explore alternative approaches to portfolio construction that go beyond conventional mean-variance analysis (Laubsch & Ulmer 1999). Scenario-based portfolio stress testing is one such methodology. It attempts to estimate the impact of extreme events and design protection against them (Basel Committee on Banking Supervision 2009; Berkowitz 2000). In applications of stress testing, one starts by identifying a relevant scenario, eg, an oil crisis or a Standard & Poor's 500 drop. The next step is modelling that event and valuing the expected losses for individual securities and the portfolio as a whole. This framework thus enables us to analyse portfolio performance subject to a concrete market event, which is making its use increasingly widespread.

Most commonly, stress testing is used to estimate the impact of adverse events on investment portfolios. In this article, we focus on how to control potential event-driven losses and risks by using a constraint-based framework within the conventional mean-variance portfolio construction process. In addition, we show how to address practical situations where stress events cannot be fully specified, and where only a select range of possible values can be assigned to model parameters. This additional complexity has not previously been considered. We introduce a robust, constrained optimisation methodology to solve for optimal stress protection in the presence of uncertainty. As a result, optimal stress hedging can now be applied more broadly to inform portfolio construction decisions. This approach can also be extended beyond the stress hedging of investment portfolios. Some examples of additional applications are a performance analysis of a loan portfolio by stress testing default correlations (Saunders & Allen 2010) and a scenario analysis of a bank's investment portfolio in the presence of common regulatory guidelines imposed as constraints in portfolio construction.

The organisation of this article is as follows. Below, we describe how to integrate optimal hedging into portfolio construction, both when a stress test scenario is fully defined and in the presence of uncertainty. Following this, our proposed methodology is illustrated, using an oil crisis scenario as an example.

Portfolio construction with stress exposure control

In portfolio construction, we can accommodate stress test exposure control by formulating it as a constrained portfolio optimisation. First, consider the case in which a scenario is described by the potential losses each asset can exhibit. In order to control aggregated portfolio loss, we extend the conventional mean-variance framework by adding a linear constraint on total or active exposure to the selected stress scenario (see, for example, Cuffe & Goldberg 2012; Ruban *et al* 2010). In such a case, the resulting portfolio construction methodology has the following form:

$$\max_{\substack{h \\ \text{given all constraints}}} [\alpha^{\mathrm{T}}(h-h_{\mathrm{b}}) - \gamma(h-h_{\mathrm{b}})^{\mathrm{T}}V(h-h_{\mathrm{b}})] \tag{1}$$
$$h \geqslant 0, \quad h^{\mathrm{T}}L \leqslant \varepsilon_1, \quad (h-h_{\mathrm{b}})^{\mathrm{T}}L \leqslant \varepsilon_2$$

where α denotes asset expected returns; V denotes the risk model; h denotes portfolio holdings; h_{b} denotes benchmark holdings; L denotes stress test loss estimates (defined to have positive values); ε_1 denotes stress test loss threshold (defined to be positive); and ε_2 denotes active stress test loss threshold (defined to be positive).

Alternatively, an investor can decide to control both stress profit and loss (P&L) and stress risk simultaneously, which results in the following optimisation formulation:

$$\max_{\substack{h \\ \text{given all constraints}}} [\alpha^{\mathrm{T}}(h-h_{\mathrm{b}}) - \gamma(h-h_{\mathrm{b}})^{\mathrm{T}}V(h-h_{\mathrm{b}})] \tag{2}$$
$$h \geqslant 0, \quad h^{\mathrm{T}}L \leqslant \varepsilon_1, \quad (h-h_{\mathrm{b}})^{\mathrm{T}}L \leqslant \varepsilon_2,$$
$$(h)^{\mathrm{T}}V_{\mathrm{S}}h \leqslant \delta_1, \quad (h-h_{\mathrm{b}})^{\mathrm{T}}V_{\mathrm{S}}(h-h_{\mathrm{b}}) \leqslant \delta_2$$

where V_{S} denotes the stress test covariance matrix, δ_1 denotes the stress test risk threshold and δ_2 denotes the active stress test risk threshold.

However, the benefit of stress exposure control does not come for free. As always, the introduction of additional constraints will come at a cost, so one should consider using a cost-benefit analysis to assess when hedging stress exposure is indeed beneficial. Multiple metrics can be utilised. For example, one can measure the cost of hedging by the amount of reduction in alpha exposure (see, for example, Ruban *et al* 2010). Alternatively, the amount of change in investor holdings can be considered as the cost of hedging. It can also be measured by either turnover or transaction costs. However, transaction cost models may not always be available, and the turnover metric, while similar to proportional transaction costs, is not efficient in asset differentiation. For these reasons, in this article we adopt a different metric. It is based on risk-scaled deviation to the current holdings. In practical terms, this approach will help to assess whether sufficient reduction in stress P&L and/or risk can be achieved without making significant changes to current investor exposures and holdings. Under this metric, the portfolio construction problem with loss and risk control in a stress scenario can be written as:

$$\min_{\substack{h \\ \text{given all constraints}}} [(h-h_{\mathrm{p}})^{\mathrm{T}}V(h-h_{\mathrm{p}})] \tag{3}$$
$$h \geqslant 0, \quad h^{\mathrm{T}}L \leqslant \varepsilon, \quad h^{\mathrm{T}}V_{\mathrm{S}}h \leqslant \delta$$

where h_p denotes current portfolio holdings, ε denotes stress test loss threshold and δ denotes stress test risk threshold.

The advantage of this measure is that it can be employed when there is no explicit information on a manager's alpha views. As a result, this approach can also be implemented for non-optimised portfolios. In addition, the methodology described by (3) is quite flexible. It can accommodate controls on alpha exposures and portfolio risk, factor in exposure bounds and incorporate bounds on total and/or active exposures to stress test scenarios.

However, there is one important limitation to all the stress test hedging solutions discussed so far in (1)–(3). These methods can be put to work only when all model parameters that describe the scenario are available, and this may not always be the case. What if an investor is not certain they can accurately estimate asset losses or a risk model describing a stress test? One simple reason for this may be the rarity of such events. Another may be that the investor tries to avoid replicating history and just assumes some of the information about their chosen stress scenario. As a result, the only information available may be the ranges of available values, or perhaps the signs for some parameters, such as correlations.

The presence of uncertainty in scenario parameters creates a problem that has not been addressed before. In this article, we tackle this complexity by using a robust optimisation approach. Our proposed technique enables one to control the worst-case outcome under stress conditions, and provides a practical portfolio construction solution.

Before proceeding with our proposed robust optimisation setup and solution, let us consider some different types of input uncertainty, previously discussed in the context of portfolio construction. Box uncertainty assumes the parameter lies within a range around the point estimate for each component (Tütüncü 2004). Ellipsoid uncertainty sets define a surface around the point estimate (centre), with axes (parameters) determining the size of the set (Goldfarb & Iyengar 2003). Both approaches can be applied to handle incomplete information in linear terms, such as stress P&L. A more complicated case is how to handle uncertainty in a stress scenario risk model. In particular, what should be done when a user can specify only the signs for correlations? While seemingly more complex, the last example can be explicitly formulated as a special case of box-type uncertainty (see, for example, Lobo & Boyd 2000). In that paper, uncertainty was present in the contemporaneous risk model used in portfolio construction. Here, we allow for uncertainty in the description of the stress scenario, but we assume the risk model corresponding to current market conditions is fully specified.

In this article, we aim to consider a rather general case of input uncertainty. Our assumption is that both the stress test covariance matrix (ie, the covariance matrix associated with the stress scenario) and the stress test loss estimates are not explicitly known. The only available information is that they are subject to a box-type uncertainty with lower and upper bounds $(\underline{C}, \bar{C})$ and $(\underline{L}, \bar{L})$, respectively. In that case, portfolio construction with stress exposure hedging can be formulated as a worst-case scenario optimisation:

$$\min_{\substack{h \\ \text{given all constraints}}} \Big[\max_{V_S}[(h-h_p)^T V(h-h_p) + \theta h^T V_S h]\Big] \tag{4}$$

$$h \geqslant 0, \quad h^T L \leqslant \varepsilon, \quad \underline{C} \leqslant V_S \leqslant \bar{C}, \quad \underline{L} \leqslant L \leqslant \bar{L}$$

where θ controls the aversion to uncertain stress scenario risk. In practice, bounds on losses can be explicitly selected based on scenario expectations, while bounds on the covariance matrix can be set by making assumptions about volatilities and correlations, which will translate into bounds on covariance. For example, when an investor wants to account for positive correlations between assets without specifying an explicit value (the bound on that correlation is particularly simple), it is between zero and one.

Using the robust optimisation techniques outlined in Lobo & Boyd (2000) and Ghaoui *et al* (2003), this mini-max problem can be cast into a convex optimisation form, which can be solved using any numerical package for semi-definite convex programming (see appendix A). We use CVX in this study (Grant & Boyd 2014).

The following illustrates an application of the proposed portfolio construction framework for stress scenario hedging, both in the case of a fully defined scenario and in the case of uncertainty.

Portfolio construction example with oil crisis scenario

To illustrate our proposed portfolio construction methodology for scenario exposure hedging, we use an oil crisis stress test as an example.[1] This 'crisis' is defined by the realisation of a $\sim 12\%$ monthly loss, as projected from daily performance in the value of crude oil futures (Nymex). We consider the 10 most recent years of data; based on historical events, the probability of oil prices plunging more than this scenario is $\sim 1\%$.

To construct deterministic stress test P&L and a covariance matrix, we employ a commonly used scenario-weighting approach similar to that described in Silva & Ural (2011) and Ruban *et al* (2010). This methodology relies on scenario specification using factor returns; for each historical observation, the distance to the selected stress scenario is computed based on the difference in risk-adjusted returns. These distances then specify the weighting scheme used to estimate the scenario model for the covariance matrix, which, in turn, is utilised to compute asset level P&L by propagating scenario shocks to other factors. In this article, we are focused on illustrating our proposed methodology for optimal hedging, so we select a single control variable – an oil price shock – and do not distinguish, for example, whether the driver for the scenario can be further classified as a supply- or demand-driven event (Kilian & Park 2009).

Also for illustration, let us assume a long-only multi-asset investment fund that has established equal-weighted positions in 10 different exchange-traded funds. Table A contains strategy correlations under normal market conditions (see the upper triangle) and current holdings. The initial portfolio allocation decision is made given the current correlation/volatility structure in the market. Alpha expectations and other portfolio manager conditions result in an *ex ante* portfolio risk of 10.7% per year. Table A also provides the results for each strategy in crisis mode. The lower triangle of the correlation matrix contains stressed values.

First, consider a deterministic case in which we only hedge against stress P&L. The resulting portfolio construction approach is described by (3), with only a linear constraint providing the upper bound for possible scenario total loss. With fully defined stress scenario parameters, we can easily compute portfolio P&L and show that, with current holdings, the fund loss will be 163 basis points if the oil crisis scenario materialises. Figure 1 shows how portfolio risk changes when potential stress test loss is reduced from 163bp to 80bp. Interestingly, our results indicate that stress protection results in lower overall portfolio risk. The green line shows that

[1] *Any of the approaches described by (1)–(3) can be further extended to simultaneously control exposure to multiple stress events.*

A. Inputs for oil crisis stress test scenario

Strategy ID	Allocation (%)	Volatility (%)	Oil crisis volatility (%)*	Oil crisis P&L (bp)	Asset correlations (current and oil crisis)**									
Strategy 1	10	16.0	19.5	−280	1.00	−0.12	0.89	0.84	0.62	−0.09	0.90	0.96	0.95	0.82
Strategy 2	10	3.0	4.1	36	−0.22	1.00	−0.06	0.08	−0.11	0.77	−0.24	−0.18	−0.18	0.15
Strategy 3	10	22.1	24.8	−329	0.87	−0.12	1.00	0.84	0.62	−0.10	0.75	0.80	0.76	0.66
Strategy 4	10	5.9	9.5	−71	0.55	0.19	0.65	1.00	0.69	0.03	0.71	0.78	0.75	0.68
Strategy 5	10	7.5	6.5	−50	0.60	0.04	0.59	0.62	1.00	−0.33	0.54	0.60	0.54	0.42
Strategy 6	10	2.3	4.5	33	−0.12	0.58	−0.07	0.19	0.19	1.00	−0.13	−0.11	−0.08	0.26
Strategy 7	10	16.2	20.4	−269	0.93	−0.24	0.74	0.49	0.57	−0.11	1.00	0.96	0.94	0.82
Strategy 8	10	17.2	20.2	−272	0.96	−0.26	0.74	0.47	0.55	−0.12	0.97	1.00	0.96	0.84
Strategy 9	10	16.6	19.4	−229	0.95	−0.25	0.73	0.46	0.54	−0.11	0.97	0.98	1.00	0.84
Strategy 10	10	16.0	20.7	−201	0.84	−0.10	0.65	0.51	0.52	−0.02	0.89	0.88	0.88	1.00

*Lowest volatility under oil crisis stress test conditions is set to be 90% of stress test risk volatility level, while highest stress test volatility is set to be 110% of stress test risk volatility for each strategy. **Upper triangle shows the strategy correlations with current risk model, while lower triangle shows the scenario-weighted correlations under oil crisis conditions. **Lower bounds on correlation uncertainty under oil crisis are calculated by subtracting 10% (50%) of absolute value of correlation oil crisis correlation if oil crisis scenario correlation is higher (lower) than current correlation for each entry. **Upper bounds on correlation uncertainty under oil crisis are calculated by adding 50% (10%) of absolute value of correlation oil crisis correlation if oil crisis scenario correlation is higher (lower) than current correlation for each entry

1 The current risk and deviation from initial holdings across various stress test P&L thresholds for a stress test P&L reduction model

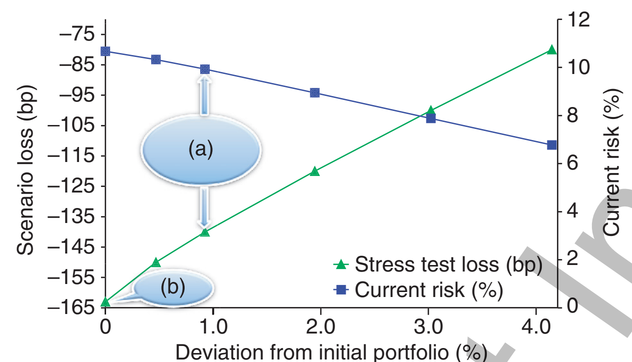

(a) Portfolio A: 23bp loss reduction at a cost of 92bp deviation from initial portfolio. (b) Initial portfolio

2 The current risk and stress test risk when stress test risk is controlled with a constraint in the stress test risk reduction model, achieving stress test P&Ls of −150bp and −140bp

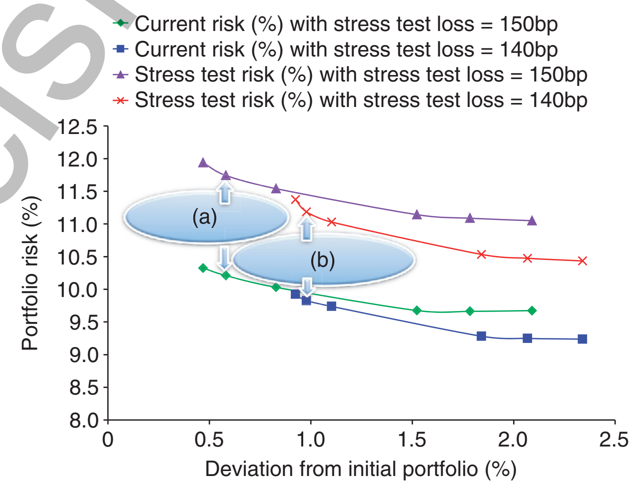

(a) Portfolio B: 13bp less loss, 87bp less stress test risk. (b) Portfolio C: 23bp less loss, 142bp less stress test risk

the risk of portfolio A decreases by 23bp if the impact of the oil crisis scenario is hedged to 140bp. The purple line shows that the cost of hedging, as measured by portfolio tracking error with respect to current holdings, increases by 92bp to achieve the hedged portfolio A. A portfolio manager needs to balance the cost of portfolio adjustments against the benefit of hedging the stress test scenario to select the appropriate reallocation solution.

Alternatively, an investor can control simultaneously for stress P&L and stress risk (see (3)). The constraint on stress test risk can be equivalently modelled as a penalty term in the objective function. Optimisations at different levels of the stress-risk aversion parameter θ will result in efficient frontiers, as shown in figure 2. On this plot, we present risk profiles of hedged portfolios against the deviation from initial holdings for varying levels of stress test risk.

The purple and red lines in figure 2 show the stress test risks of different hedged solutions that achieve 150bp and 140bp losses under the oil crisis scenario. When we compare these with the green and blue lines, we can see that scenario-weighted stress test risk is 14% more than the current risk on average. Portfolio B shows a hedged solution that reduces oil crisis risk by 142bp and oil crisis loss by 40bp, with a 1% deviation from the current holdings. Compared with portfolio C, portfolio B achieves a 10bp greater reduction in ‘oil crisis’, with 55bp less stress test risk, 38bp less current risk and 39bp more deviation from initial holdings at the same level of stress test risk aversion.

While actual numbers depend on model details, the overall shape of these efficient frontiers can be explained qualitatively. Indeed, the simultaneous improvement in stress P&L and stress risk exhibited by portfolio B versus portfolio C is in line with the expectation that stress test risk and stress test loss reduction are positively correlated. The hedging opportunities increase with a higher loss reduction level (blue line versus green line), leading to higher benefits of risk reduction per one unit of deviation from initial holdings. However, there is a limit to risk reduction for a given level of stress test loss. The curvature of the blue line shows the stress test reduces at diminishing rates with further deviation from initial holdings. In

B. The upper and lower bounds on volatility and correlations for the oil crisis scenario

Strategy ID	Lower volatility bound (%)	Upper volatility bound (%)	Asset correlations (upper & lower bound for oil crisis)***									
Strategy 1	17.6	23.4	1.00	−0.20	0.96	0.61	0.66	−0.11	1.00	1.00	1.00	1.00
Strategy 2	3.6	4.9	−0.33	1.00	−0.11	0.29	0.06	0.64	−0.12	−0.23	−0.23	−0.09
Strategy 3	22.3	29.8	0.44	−0.18	1.00	0.72	0.65	−0.04	0.81	0.81	0.80	0.72
Strategy 4	8.6	11.4	0.28	0.17	0.33	1.00	0.68	0.29	0.54	0.52	0.51	0.56
Strategy 5	5.8	7.8	0.30	0.04	0.30	0.31	1.00	0.29	0.86	0.61	0.81	0.78
Strategy 6	4.0	5.3	−0.18	0.29	−0.08	0.17	0.17	1.00	−0.06	−0.11	−0.10	−0.02
Strategy 7	18.4	24.5	0.84	−0.26	0.37	0.25	0.51	−0.12	1.00	1.00	1.00	1.00
Strategy 8	18.2	24.2	0.86	−0.39	0.37	0.24	0.28	−0.18	0.87	1.00	1.00	1.00
Strategy 9	17.4	23.2	0.86	−0.38	0.37	0.23	0.49	−0.17	0.87	0.88	1.00	1.00
Strategy 10	18.6	24.8	0.76	−0.15	0.33	0.26	0.47	−0.03	0.80	0.79	0.79	1.00

***Upper triangle shows the upper bound on strategy correlations under oil crisis scenario, while lower triangle shows the lower bound on correlations for oil crisis scenario

practice, computing these efficient hedging frontiers will inform investors on the desired level of hedging.

To illustrate our approach for a scenario with uncertainty, we decided to use point estimates for stress P&L, while allowing box-type constraints for elements of the stress covariance matrix. Here, we use the following rules to specify the range for correlations and volatilities, so they encompass observed variability across all periods identified by the stress conditions.[2] The correlation upper bound is defined as the stress-based estimated value, increased (decreased) by 50% of the absolute value of stress regime correlation if the stress regime correlation is higher (lower) than the normal regime correlation. The lower bound is defined as the estimated stressed value, reduced (increased) by 10% of the absolute value of the stress regime correlation if the stress regime correlation is higher (lower) than the normal regime correlation. This asymmetry reflects the expectation that correlations tend to increase during crisis periods. In addition, for each strategy we allowed volatility to be within 90–120% of the oil crisis values. The resulting risk model (table B) combines uncertainty coming from both sources.

The effect of robust optimisation was measured by changing the stress risk aversion parameter to determine the impact of the worst-case scenario on portfolio allocations. This is similar to our analysis for a fully defined stress risk model, which allows us to compare both approaches.

In figure 3, we depict efficient stress hedging frontiers with box-type uncertainty in a stress test risk model. This figure shows how the current risk (green line), the worst-case risk (blue line) and the deviation from the initial portfolio (purple line) change with the increasing impact of stress test risk uncertainty (horizontal axis), while the expected stress test scenario loss is kept constant at 140bp. As expected, worst-case risk values can be significantly higher than the current portfolio risk estimate. In fact, in this example, the worst-case volatility associated with the crisis scenario is almost 50% higher than the value based on a regular model.

Stress test hedging with robust optimisation provides both a worst-case scenario covariance matrix and optimal holdings. When the upper bounds on the covariance matrix yield a positive semi-definite matrix, we can simply use the upper bound as the worst-case risk covariance. However, when the upper bounds on the covariance matrix are not positive semi-definite, the robust optimisation approach guarantees the worst-case scenario risk model is a valid risk model. Asset correlations do not move in

3 How portfolio risk and deviation from initial holdings change when stress test risk trade-off is varied for the stress test hedging with the robust optimisation model for a fixed stress test loss threshold of −140bp

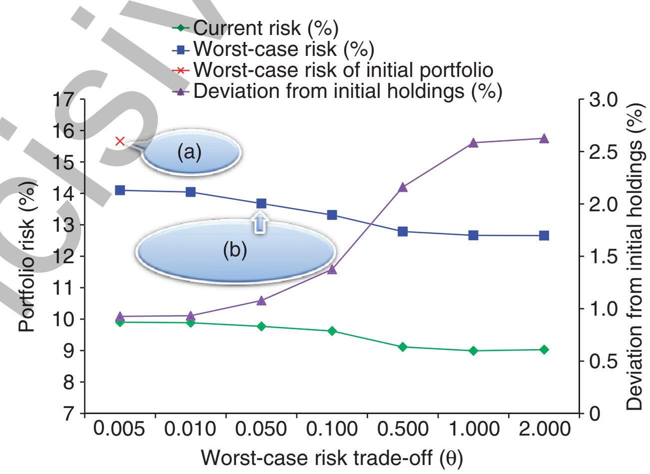

(a) Initial portfolio. (b) Portfolio D: 23bp less loss, 198bp less worst-case risk

C. Upper bounds on stress test covariance matrices are not always achieved with the stress test hedging with robust optimisation model for portfolio D

	Strategy 1	Strategy 2	Strategy 4	Strategy 9
Strategy 1	1	−0.20	0.61	1.00
Strategy 2	−0.23	1	0.29	−0.23
Strategy 4	0.61	0.28	1	0.51
Strategy 9	0.98	−0.23	0.51	1

Red indicates upper bound. Green indicates upper bound achieved. Blue indicates upper bound not achieved

the same direction under the oil crisis scenario, and the upper bounds on asset correlations are not always achieved. Some asset correlations under the worst-case scenario for portfolio D are shown in table C. For example, the correlation between strategy 1 and strategy 9 cannot be 1.0 in the worst-case scenario.

To illustrate the efficiency of robust optimisation in reducing worst-case scenario risk, we compare two portfolios that have the same level of loss reduction. Portfolio A is based on loss hedging in a deterministic scenario (see figure 1). Portfolio D achieves the same level of loss protection and is

[2] *The number of observations was not sufficient to model and explore variability in risk model uncertainty across different stress periods.*

a solution for robust optimisation for scenario risk control in the presence of uncertainty. We can estimate the worst-case risk for both portfolios; as expected, including risk control results in a lower value of worst-case volatility.

Summary

Stress test analysis provides an efficient framework with which to identify the potential impact of a market event on an investment portfolio. In this article, we describe how to control that negative impact by constraining the loss and/or risk associated with the stress scenario in portfolio construction. Furthermore, we propose a novel, robust constrained optimisation methodology that directly addresses the challenges posed by uncertainty when modelling stress test scenarios. Using the computational example of an oil crisis, we illustrate the application of our proposed methodologies and the effectiveness of the resulting solutions.

Finally, the stress test hedging framework presented is quite flexible. It can be extended to multiple scenarios by simultaneously controlling the different drivers of risks. In addition, the proposed robust optimisation approach can not only handle uncertainty in asset-by-asset risk models but also be applied to partially defined factor-based structural risk models.

Appendix A: robust optimisation approach for stress test hedging

We consider the same problem formulation described in (4). The box-type uncertainty in stress covariances is set by its lower and upper bounds $(\underline{V}, \bar{V})$, which can be driven by, for example, uncertainty in correlations, while stress test loss is represented with the following set:

$$L = \{L\colon L_0 + \gamma,\ |\gamma_i| \leqslant \mu_i,\ i = 1, \dots, n\}$$

where L_0 shows the point estimate for scenario loss and γ is the uncertainty around each scenario loss estimate.

After using the techniques outlined in Lobo & Boyd (2000) and Goldfarb & Iyengar (2003), this problem can be cast into a convex robust optimisation model as follows:

$$\min_{h, h^+, h^-, \underline{Q}, \bar{Q}} [(h - h_\mathrm{p})^\mathrm{T} V (h - h_\mathrm{p}) + \theta(\langle \bar{Q}, \bar{V} \rangle - \langle \underline{Q}, \underline{V} \rangle)]$$

$$h^\mathrm{T} L_0 + \mu^\mathrm{T}(h^+ + h^-) \leqslant \varepsilon, \quad h = h^+ - h^-, \quad h^+ \geqslant 0, \quad h^- \geqslant 0,$$

$$\underline{Q} \succeq 0, \quad \bar{Q} \succeq 0, \quad M = \begin{bmatrix} \bar{Q} - \underline{Q} & h \\ h^\mathrm{T} & 1 \end{bmatrix} \succeq 0$$

The robust model includes auxiliary variables $(h_+, h^-, \underline{Q}, \bar{Q})$ in order to choose a worst-case risk model from the covariance uncertainty set. Here, $\langle A, B \rangle$ represents the trace of the matrix product AB. The trade-off between the deviation from the initial holdings and the worst-case stress test risk is controlled by θ. The inequality constraints ($\succeq$) on matrices $\underline{V}, \bar{V}, \underline{Q}, \bar{Q}$ define the positive semi-definite structure of these matrices. To illustrate this approach, we include only the uncertainty in the stress scenario risk model described above. The resulting convex semi-definite optimisation problem was solved using Grant & Boyd (2014). ■

Mehmet Bilgili is a vice president and Maurizio Ferconi and Alex Ulitsky are managing directors and at BlackRock in San Francisco. **The authors report no conflicts of interest. The authors alone are responsible for the content and writing of this article. The views expressed here are those of the authors alone. This material is not the property of BlackRock and is not representative of the views of BlackRock, its officers or directors. This note is intended to stimulate further research and is not a recommendation to trade particular securities or of any investment strategy.** **The authors thank Stephen Boyd, Ronald Ratcliffe, Vivek Gupta and anonymous referees for their valuable input and comments.**
**Email: Mehmet.Bilgili@blackrock.com,
Maurizio.Ferconi@blackrock.com,
Alex.Ulitsky@blackrock.com.**

REFERENCES

Basel Committee on Banking Supervision, 2009
Principles for sound stress testing practices and supervision
Report, Bank for International Settlements

Berkowitz J, 2000
A coherent framework for stress testing
Journal of Risk 2(2), pages 5–15

Cuffe SL and LR Goldberg, 2012
Allocating assets in climates of extreme risk: a new paradigm for stress testing portfolios
Financial Analysts Journal 68(2), pages 85–108

Ghaoui LE, M Oks and F Oustry, 2003
Worst-case value-at-risk and robust portfolio optimization: a conic programming approach.
Operations Research 51(4), pages 543–556

Goldfarb D and G Iyengar, 2003
Robust portfolio selection problems
Mathematics of Operations Research 28(1), pages 1–38

Grant M and S Boyd, 2014
CVX: Matlab software for disciplined convex programming, version 2.1
Available at *http://cvxr.com/cvx*

Kilian L and C Park, 2009
The impact of oil price shocks on the US stock market
International Economic Review 50(4), pages 1267–1287

Laubsch AJ and A Ulmer, 1999
Risk Management: A Practical Guide
RiskMetrics Group

Lobo MS and S Boyd, 2000
The worst-case risk of a portfolio
Unpublished Manuscript, available at http://bit.ly/2qV7x5c

Ruban OA, D Melas and MSCI Inc, 2010
Stress testing in the investment process
Research Paper 2010-28, August, MSCI Barra

Saunders A and L Allen, 2010
Credit Risk Management In and Out of the Financial Crisis: New Approaches to Value at Risk and Other Paradigms, volume 528
John Wiley & Sons

Silva A and C Ural, 2011
Stress testing of portfolios
Report, November, Barclays Capital Index, Portfolio and Risk Solutions

Tütüncü RH, 2004
Robust asset allocation
Annals of Operations Research 132(1-4), pages 157–187